Cooking Lightweight XML Query Processor with Binary Joins and Comparing it with Holistic Joins: Technical Report

Petr Lukáš
petr.lukas@vsb.cz

Radim Bača
radim.baca@vsb.cz

Michal Krátký
michal.kratky@vsb.cz

VŠB – Technical University of Ostrava
Department of Computer Science, Database Research Group
17.listopadu 15, Ostrava, Czech republic

arXiv:1703.09539v2 [cs.DB] 26 Apr 2017

ABSTRACT

XML queries can be modeled by twig pattern queries (TPQs) specifying predicates on XML nodes and XPath relationships satisfied between them. A lot of TPQ types have been proposed; this paper takes into account a TPQ model extended by a specification of output and non-output query nodes since it complies with the XQuery semantics and, in many cases, it leads to a more efficient query processing. In general, there are two approaches to process the TPQ: holistic joins and binary joins. The holistic joins have been developed as a generalization of the binary joins, and they have been considered as a state-of-the-art TPQ processing method. Surprisingly, a thorough analytical and experimental comparison is still missing despite an enormous research effort in this area. In this paper, we try to fill this gap; we analytically and experimentally show that the binary joins used in a fully-pipelined plan (i.e., the plan where each join operation does not wait for the complete result of the previous operation) can often overcome the holistic joins even without any cost-based optimizer, especially for TPQs with a higher ratio of non-output query nodes. The main contributions of this paper can be summarized as follows: (i) we introduce several improvements of existing binary join approaches allowing to build a fully-pipelined plan for a TPQ considering non-output query nodes, (ii) we prove that for a certain class of TPQs such a plan has the linear time complexity with respect to the size of the input and output as well as the linear space complexity with respect to the XML document depth (i.e., the same complexity as the holistic join approaches), (iii) we show that our improved binary join approach overcomes the holistic join approaches in many situations.

1. INTRODUCTION

Searching for occurrences of a twig pattern in an XML tree is often considered as a core problem of the structural XML query processing. An enormous research effort has been devoted to this problem and we can identify two major approaches: (1) binary join approaches [1, 2, 20], and (2) holistic join approaches [7, 8, 9, 15, 16, 12, 3]. *The binary join approaches* are often considered to be very flexible in terms of their simple integration into any XML algebra [5, 13] and the support of all XPath axes. The main disadvantage is the possible large intermediate result containing many data nodes irrelevant to a TPQ [7]. *The holistic join approaches* have been introduced to overcome this problem. They prove the linear time complexity with respect to the size of the output and input and the linear space complexity with respect to the XML document depth for a certain class of TPQs. This feature leads to their wide popularity and the appearance of many various extensions and improvements.

Despite the theoretical advantages of holistic join approaches, a thorough experimental comparison with binary join approaches is still missing in the literature. The only experimental comparison can be found in [7]; however, it was an early stage of the TPQ processing and many advances have occurred since then. One aim of this article is to compare both types of the approaches and to dispel myths about them.

Since a TPQ is possibly processed with a lot of binary joins, the building of an appropriate query plan is a crucial problem of the binary join approaches. A *fully-pipelined* (FP) query processing plan is important for the binary join approaches [26], since a join operation of such a plan does not wait for the complete result of the previous join operation, i.e., the join operation result is not materialized. However, it is possible to utilize an FP plan only when the previous operation provides its result ordered as the next operation requires. In the opposite case, an expensive sort operation has to be performed over the complete result and the result cannot be therefore consecutively passed to the next operation.

Following the semantics of XQuery, we can identify two types of query nodes: (i) *output query nodes* corresponding to 'for' clauses and (ii) *non-output query nodes* corresponding to all other query nodes. In [26], the authors describe how to build an FP plan for a TPQ where all query nodes are output and the result is sorted according to an arbitrary

query node. However, there are two issues that should be taken into account when XQuery is considered:

1. Neglecting the semantics of non-output query nodes has a critical impact on the complexity of the TPQ processing [3, 25].

2. Sorting the TPQ result by one query node is not sufficient for XQuery queries where the result have to be lexicographically sorted by all output query nodes according to the order of the 'for' clauses in the query.

In this article, we combine ideas proposed in [26, 2, 20] and add certain modifications of binary join approaches; in this way we solve these issues and show how to build an FP plan for such a TPQ. Moreover, we show that it is not necessary to utilize a cost-based optimizer to build an FP plan. Let us note that the holistic join approaches also do not often utilize a cost-based optimizer. As a result, we simply consider a cost-based optimization to be orthogonal to both types of approaches [11, 4] and we completely avoid it in our work. In this way, we can compare the binary join approach with holistic join approaches. As a bonus, we can prove that for a certain class of TPQs our binary join approach has the linear time complexity with respect to the size of the input and output and the linear space complexity with respect to the depth of the XML document. To the best of our knowledge, this is the first work which proves such a complexity of the binary join approaches.

The main contributions of this paper are summarized as follows:

1. We introduce our binary join approach based on FP plans which is an extension of the existing ones. It builds FP plans for a TPQ considering non-output query nodes without a cost-based optimizer.

2. We prove time and space complexities of our binary join approach for a certain class of TPQs. We show that the complexity of the binary join approach bears many similarities to the complexity of holistic join approaches. The only difference is that the complexity of holistic approaches is not influenced by the number of non-output query nodes.

3. An experimental comparison of both approaches. We clearly describe use-cases in which each of them excels. We show that our binary join approach is more likely to be advantageous over the holistic join approach for TPQs with a higher ratio of non-output query nodes, and this experimental observation is in the compliance with the complexity analysis.

Section 2 depicts the TPQ processing problem and Section 3 gives an overview of techniques that are used to solve this problem. Section 4 includes a thorough description of the binary join algorithms and Section 5 describes how to build an FP plan based on these algorithms. Section 6 conducts a thorough complexity analysis of our binary join approach. Section 7 experimentally compares both approaches to distinguish their weak and strong aspects.

2. PRELIMINARIES

Let us first define some basic terms we use in this paper.

2.1 Document and Query Model

An *XML document* (or simply a *document*) is an ordered rooted tree. The nodes in the tree are called *data nodes*. Each data node is labeled with a *tag*. A sample XML document is depicted in Figure 1a.

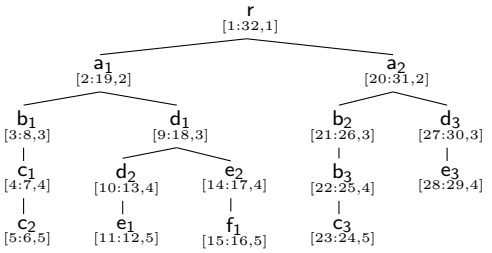

(a) XML Document

```
for $i in //r/a
for $j in $i/b//c
for $k in $i//d[./e and .//f]
return ($i, $j, $k)
```

(b) Query **Q1**

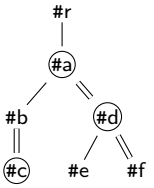

(c) TPQ of Q1

Figure 1: A sample document, an XQuery query, and a TPQ

A *twig pattern query (TPQ)* is also an ordered rooted tree. The nodes in the tree are called *query nodes*. Each query node is labeled with: (1) an expression $\#\theta$ specifying all data nodes with the tag θ and (2) bit information whether the query node is an *output query node* or a *non-output query node*. The edges represent XPath relationships; we consider only *ancestor-descendant* (AD) and *parent-child* (PC) relationships in this paper. Since a TPQ models an XML query, we use the terms 'query' and 'TPQ' interchangeably. An example of a TPQ is illustrated in Figure 1c. Single and double lines represent PC and AD edges, respectively, the output query nodes are circled. The TPQ corresponds to the sample XQuery query Q1 in Figure 1b.

Let us note that there are other query models covering the semantics of the TPQ with non-output query nodes in the literature. For example, a generalized twig pattern (GTP) is introduced in [10] and [25] defines a TP+Output model. We do not utilize these richer models since our focus is on the semantics related to the 'for' clause (i.e., output query nodes) due to its importance for the query processing efficiency. However, in Appendix D, we show that the GTP model covering the semantics of 'let' clauses and ('or' and 'not') logical connectives can be easily integrated into our binary join approach and it is actually implemented in our prototype.

In this article, we define some specific parts of a query. A *core query node* is an output query node or any query node lying on a path between two output query nodes. A *query core* is a query subtree formed of all core query nodes. A *predicate* Π_q for a core query node q (called the *output predicate node*) is a query subtree formed of q and all non-output query nodes connected with q by edges not lying on a path between two output query nodes. In Section 6, we show that the use of different join algorithms for the query core and for the predicates leads to a more efficient query processing.

EXAMPLE 2.1. *Let us consider the query Q1 in Figure 1c. There are 4 core query nodes:* #a, #b, #c, *and* #d *(see Figure 2a). The predicates for* #a *and* #d *are depicted in Figure 2b and 2c, respectively. The predicates for* #b *and* #c *are only the core query nodes themselves.*

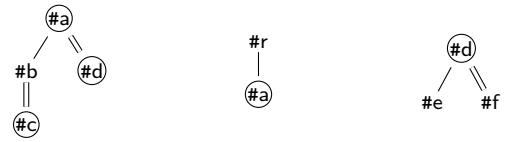

(a) Query core (b) Predicate $\Pi_{\#a}$ (c) Predicate $\Pi_{\#d}$

Figure 2: Subtrees of the query Q1

2.2 Problem Statement

Having a document D and a query Q with n query nodes $q_1, \cdots, q_n$. A *complete match* is an n-dimensional tuple $[d_1, \cdots, d_n]$, where $d_1, \cdots, d_n$ are data nodes from D specified by the query nodes $q_1, \cdots, q_n$ such that the relationships between $d_1, \cdots, d_n$ satisfy the AD or PC relationships defined by edges between $q_1, \cdots, q_n$. An *output match* is a projection of a complete match such that data nodes corresponding to non-output query nodes are omitted. The result of the query Q for the document D is an ordered set of all output matches. The output matches are lexicographically sorted by the document order of data nodes corresponding to the output query nodes in a pre-order traversal.

EXAMPLE 2.2. *The result of Q1 in Figure 1c for the document in Figure 1a is shown in Table 1. We see that the output query nodes are ordered in a pre-order traversal.*

#a	#c	#d
a_1	c_1	d_1
a_1	c_2	d_1

Table 1: The result of the query Q1

Let us note that there exist XQuery queries that cannot be modeled by a TPQ such that the result is ordered correctly. We assume that these queries are rather impractical and rare and we do not consider them in this paper. An example of such a query is shown in Appendix A.

2.3 Labeling Scheme

Each data node in a document is labeled with a unique *node label*[1, 7, 3]. The node label allows us to resolve XPath relationships between data nodes. In this paper, we use the containment labeling scheme [28], where a node label is a 3-tuple $[left : right, level]$. Using the containment scheme, a data node a is an ancestor of a data node d if the interval $(a.left, a.right)$ contains the interval $(d.left, d.right)$. Similarly, a data node p is the parent of a data node c if p is an ancestor of c and $p.level + 1 = c.level$. Let us note that any labeling scheme which is capable of resolving AD or PC XPath relationships (e.g., Dewey order [21]) can be used as well. In Figure 1a, we see an XML tree with node labels of the containment labeling scheme. Further in the text, we use the terms 'data node' and 'node label' interchangeably.

2.4 Query Plan

An *n-ary operator* is a mapping $T_{in_1} \times \cdots \times T_{in_n} \rightarrow T_{out}$, where $T_{in_1}, \cdots, T_{in_n}$ are input streams and T_{out} is the output stream. A *stream* T is an abstract data structure supporting the following operations:

- `advance(`T`)` – moves to the next item in T,
- `finished(`T`)` – returns a boolean value whether the whole stream T has been completely iterated,
- `current(`T`)` – returns the current item of T.

$|T|$ denotes the size of the stream T, i.e., the number of advance calls until the stream is finished.

In this paper, we combine operators working with a *table* [1, 2, 3] and operators working with a *list* [20]; both data structures implement the stream. A table is an array of tuples of data nodes lexicographically sorted by one or more columns using the document ordering, a list is the equivalent of a one-column table.

A *query plan* is an ordered rooted tree of operators where inputs of an operator are outputs of its children. A *fully-pipelined (FP) plan* [26] is a query plan where each operator o produces the output stream sorted such that the parent of o (or an end-user if o is the root) requires. The FP plans bring the following advantages: (1) an explicit sorting is not necessary, and (2) the parent operator can consecutively read the result without waiting for the complete result of the previous operator (i.e., the operators can be implemented as cursors).

2.5 Query Processing Approaches

We can distinguish two TPQ processing approaches: (i) *binary join approaches* and (ii) *holistic join approaches*. The binary join approaches utilize a query plan with a number of *binary join* operators where each join is related to one edge in the query. The holistic join approaches utilize a query plan with one operator – a *holistic join*. We define the binary and holistic join operators in the following subsections. Let us note that there are approaches combining both types of joins [23].

Both types of approaches utilize an operator *IndexScan* (IS). The $IS(q)$ operator retrieves data nodes specified by q, where q is a tag. We suppose an index data structure implementing this operator (e.g., B^+–tree).

2.5.1 Binary Join

In the literature, the binary join operator is sometimes called a structural [1, 26] or a containment join [2]. There are many types of this operator which can be categorized from various points of view. Following the specific parts of a query defined in Section 2.1, our approach uses *semi-joins* [20] to process the query predicates and *partial-joins* [2] to process the query core. The semi-joins are derived from the partial-joins and their implementation is more efficient, however, the semi-joins cannot be used to process the query core.

The *partial-join* operator [2] is defined as follows:

$$T_o = PartialJoin_{[M_a, M_d, i, j, \alpha]}(T_a, T_d),$$

where T_a and T_d are tables with m and n columns, M_a and M_d are binary *projection masks* with m and n bits specifying which columns from T_a and T_d are projected to the output, i and j are indexes of *join columns* in T_a and T_d, respectively.

α is a required XPath relationship (AD or PC). The partial-join operator requires T_a and T_d to be sorted by the join columns. T_o is the output table with o columns ($o = m' + n'$), where m' and n' are numbers of non-zero bits in M_a and M_d, respectively. T_o is computed by joining all tuples in T_a and T_d by the join columns satisfying the relationship α. T_o is lexicographically sorted by data nodes from the join columns of T_a or T_d. Thus, we have an *ancestor-sorted* or a *descendant-sorted* variant of the partial-join, respectively.

The *semi-join* operator [20] is defined as:

$$T_o = SemiJoin_{[\alpha]}(T_a, T_d),$$

where T_a and T_d are lists and α is a required XPath relationship. There are two variants of the semi-join: (i) an *ancestor-filtering* variant produces a list T_o of all data nodes $a \in T_a$ such that there is a data node $d \in T_d$ in the relationship α with a, and (ii) a *descendant-filtering* variant produces a list T_o of all data nodes $d \in T_d$ such that there is a data node $a \in T_a$ in the relationship α with d.

2.5.2 Holistic Join

A holistic join is an operator with n input lists $T_{q_1}, \ldots, T_{q_n}$ of data nodes producing a table of all output matches. There exists a number of algorithms implementing a holistic join [7, 8, 9, 15, 12, 3]. When compared to a binary join, the major difference is that a holistic join considers current items of all streams in every step of the algorithm, whereas a binary join considers only a pair of streams.

3. RELATED WORK

Various TPQ processing approaches have been developed during the last two decades [1, 7, 8, 9, 15, 16, 28, 12, 3]. Two major types of XML processing approaches are represented by the binary join approaches [1, 28, 27, 14, 2] and the holistic join approaches [7, 8, 9, 15, 16, 12, 3]. The binary join approaches can be easily integrated into a full-fledged XQuery processor in order to support all XPath axes [10, 18, 6], however, large intermediate results can be produced in certain cases (see Section 6 for details). In [1], the *StackTreeAnc* and *StackTreeDesc* algorithms are introduced; these algorithms are considered as traditional binary (structural) joins used by many approaches [2, 20, 26]. In Section 4.1, we present an extension to *StackTreeAnc* allowing us to build an FP plan for a TPQ considering the non-output query nodes semantics. In [2], the authors introduce a partial-join which can project out some columns of the input tables. Their experiments show that incorporating the partial-joins can reduce the query processing time. In [20], a structural semi-join operator is introduced; the semi-join is a special case of the partial-join presented in [2]. They present algorithms *NStack-SSJoin-Desc* and *NList-SSJoin-Anc* that implement the descendant-filtering and ancestor-filtering semi-joins, respectively. *NStack-SSJoin-Desc* is implemented without any stack and *NList-SSJoin-Anc* is implemented without self- and inherited-lists, but the stack is still necessary. In Section 4.2, we introduce a stack-less variant of the ancestor-filtering semi-join.

In [26], the authors put forward several cost-based algorithms to select a query plan that is close to the optimal query plan. Their experiments show that the FP plan provides the most efficient performance when the complete query processing time is considered. The authors also show how to build the FP plan producing a table sorted by an arbitrary query node; however, they do not consider non-output query nodes and lexicographic sorting by multiple columns which is necessary to support the XQuery semantics. In Section 5, we introduce an algorithm building the FP plans that overcomes these limitations. To the best of our knowledge, this is the first approach where the lexicographic sorting is considered without an explicit sorting of the result.

Holistic join approaches reduce the intermediate result size without any cost-based optimization and they provide precise theoretical guarantees [7]. Holistic join approaches can be also integrated into an XQuery algebra [17] and even combined with binary join approaches in the same query plan [24]. However, the number of XQuery constructs that can be processed holistically is limited, e.g., only some XPath axes are supported. There are many various holistic join algorithms that can be categorized from different perspectives. The first perspective relates to the way the stacks and an intermediate result data structure are used: (1) the top-down holistic joins [7, 15, 16] usually skip most of irrelevant nodes before they are stored on stacks, (2) the bottom-up holistic joins [8] usually skip most of irrelevant nodes when they are popped out from their stacks, and (3) the combined holistic joins [12, 3] employ advantages from both types of holistic joins. The second perspective categorizes holistic joins according to a query model used: (1) approaches that do not reflect the notation related to the output query nodes [7, 15, 16, 12], and (2) approaches that work with the output query nodes and produce a table of output matches [8, 3].

We pick the *GTPStack* holistic join [3] as a representative of the state-of-the-art holistic algorithms in our experimental comparison due to the fact that it combines the advantages of top-down and bottom-up approaches and it reflects the output query node semantics.

4. BINARY JOIN ALGORITHMS

In this section, we describe algorithms implementing the partial-join and semi-join operators with focus on several improvements enabling us to build an FP plan. A thorough description how to build an FP plan is given in Section 5.

Join type	Relation-ship	anc.-sorted anc.-filtering	desc.-sorted desc.-filtering
part.-join	AD, PC	*StackTreeAnc* [1]	*StackTreeDesc* [1]
semi-join	AD	*FilterAncAD*	*FilterDescAD*
	PC	*FilterAncPC*	*FilterDescPC*

Table 2: Summary of binary join algorithms

The partial-join algorithms *StackTreeAnc* and *StackTreeDesc* were introduced in [1] and a few modifications regarding their signature were presented in [2]. The semi-join algorithms borrow ideas from [20, 2]. In the case of the partial-joins, the (AD or PC) XPath relationship is just a parameter of the algorithms. For the semi-joins, both XPath relationships are implemented using different algorithms, since the processing of the AD relationship can be done more efficiently than the PC relationship.

4.1 Partial-join Algorithms

The partial-join algorithms *StackTreeAnc* and *StackTreeDesc* are based on merging two streams of node labels to satisfy the AD or PC relationship. Both algorithms require a stack to keep a track of all ancestors of any data node from T_a at any time of the query processing. *StackTreeAnc* also requires extra data structures to produce a sorted result. These data structures are called the self- and inherited-lists and each entry in the stack of *StackTreeAnc* is associated with one self- and one inherited-list. A detailed description of both algorithms is given in [1].

We can identify a sorting problem in the case of queries with non-output core query nodes. In Figures 3a and 3b, we see two similar queries Q2 and Q3. The only difference of Q3 is that #d is a non-output query node. Before we describe the sorting problem which appears in the case of Q3, let us show how the partial-joins are used in an FP plan to process Q2.

EXAMPLE 4.1. *The result of Q2 for the document in Figure 1a icnludes all triplets* [a, d, e] *such that* a *is an ancestor of* d *and* d *is the parent of* e. *The result is sorted lexicographically by* (a, d, e).

(a) Query **Q2** (b) Query **Q3**

Figure 3: Sample TPQs

A query plan of Q2 is depicted in Figure 4a. Since all query nodes are output, we have to use only the partial-join algorithms (see Section 5). To produce the sorted result, we need to join the query nodes in a bottom-up manner, so the last join operation produces the result sorted by #a. *We first join* #d *with* #e *to produce all pairs* [d, e] *such that* d *is the parent of* e. *We use the ancestor-sorted partial-join (i.e., StackTreeAnc), so the intermediate result (see Figure 4b) is sorted lexicographically by* (d, e). *Subsequently, we join the intermediate result with* #a *to produce all triplets* [a, d, e] *such that* a *is an ancestor of* d*; we use StackTreeAnc again to make the final result (see Figure 4c) sorted lexicographically by* (a, d, e).

Now, let us consider the query Q3 in Figure 3b, where #d is a non-output query node. We cannot use the plan in Figure 4a, since the order of the e data nodes would be corrupted because of the recursive data nodes d_1 and d_2. A naive solution is to use the plan in Figure 4a and then to sort the final result by (a, e). However, we lose the advantages of the FP plan since a sorting operation is used (see Section 2.4). To solve this problem without any additional sorting, we introduce an extension of the *StackTreeAnc* algorithm called a *secondary relationship test*: we extend the signature of the algorithm with an attribute k.

$$T_o = StackTreeAncSrt_{[M_a,M_d,i,j,\alpha,k]}(T_a, T_d),$$

where j and k refer to the *primary* and the *secondary* join columns in T_d, respectively. Let c_i be the join column of

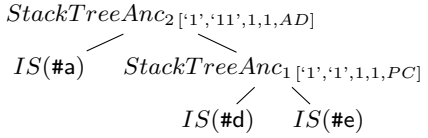

(a) A query plan of Q2

#d	#e
d_1	e_2
d_2	e_1
d_3	e_3

(b) The intermediate result – the result of $StackTreeAnc_1$

#a	#d	#e
a_1	d_1	e_2
a_1	d_2	e_1
a_2	d_3	e_3

(c) The final result – the result of $StackTreeAnc_2$

Figure 4: An example of the query processing

T_a, let c_j and c_k be the primary and the secondary join columns in T_d, respectively. The output is computed such that the input tables are joined by data nodes in c_i and c_j using the traditional *StackTreeAnc* algorithm with the AD relationship, even if α is PC. It requires the table T_d to be sorted by c_j. However, the joined tuple is output iff the relationship α is satisfied between data nodes from c_i and c_k. This is called the secondary relationship test.

EXAMPLE 4.2. *When the novel algorithm is applied, we solve the sorting problem of the query Q3 (see Figure 3b) using the plan in Figure 5. We first join* #d *with* #e *using a descendant-sorted partial-join (i.e., the StackTreeDesc algorithm), so the pairs* [d, e] *in the intermediate result are sorted lexicographically by* (e, d). *Subsequently, we use the novel StackTreeAncSrt algorithm joining the intermediate result with data nodes of* #a *to produce all triplets* [a, d, e], *such that* a *is an ancestor of* e*; using the secondary relationship test, we verify whether* a *is an ancestor of* d. *As a result, the triplets are sorted lexicographically by* (a, e, d). *To filter-out* d *data nodes, we set the first bit of M_d of the StackTreeAncSrt to zero.*

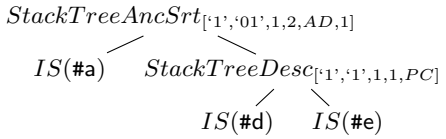

Figure 5: A plan solving the sorting problem using the novel *StackTreeAncSrt* algorithm

4.2 Semi-join Algorithms

In [2, 20], the authors show that a descendant-filtering semi-join for the AD relationship can be implemented without any stack, which leads to a more efficient query processing. In our article, we call this algorithm *FilterDescAD*. To the best of our knowledge, a stack-less ancestor-filtering semi-join for the AD relationship has not yet been presented in the literature; in this section, we present such an algorithm called *FilterAncAD*.

The descendant- and ancestor-filtering semi-joins for the PC relationship are a straightforward simplification of *StackTreeDesc* and *StackTreeAnc* [1] where only the nodes of T_d

or T_a form the output. In this paper, we call the semi-joins *FilterDescPC* and *FilterAncPC*, respectively. Unfortunately, using of a stack is unavoidable in the case of PC. In [20], a stack-less descendant-filtering semi-join *NStack-SSJoin-Desc* is proposed, but it is not clear how the algorithm deals with the PC relationship.

Algorithm 1: FilterAncAD(T_a, T_d)

```
   input  : Two lists T_a and T_d
   output: A list of data nodes in T_a having a data node
           with the AD relationship in T_d
1  while ¬ finished(T_a) ∧ ¬ finished(T_d) do
2      if current(T_a).left ≥ current(T_d).left then
3          advance(T_d);
4      else if current(T_a).right > current(T_d).right
       then
5          output current(T_a);
6          advance(T_a);
7      else
8          advance(T_a);
9
```

The novel *FilterAncAD* is shown in Algorithm 1. Lines 1 – 9 loop through all nodes in T_a and T_d. In Line 3, the stream T_d is advanced since its current data node precedes the current data node of T_a. Lines 5 – 6 cover the case when the current node of T_d is a descendant of the current node of T_a. Otherwise, the current node of T_a precedes the current node of T_d and the stream T_a needs to be advanced.

5. BUILDING FP PLAN

In this section, we show how to combine the proposed partial-join and semi-join algorithms to build an FP plan for a TPQ considering non-output query nodes. An algorithm building an FP plan is presented in [26]. However, this work does not consider non-output query nodes, which is necessary to process XQuery queries. We also show that considering the non-output query nodes means the more efficient processing of queries since semi-joins can be used instead of partial-joins (see Section 7.1). Moreover, sorting by multiple columns has to be considered to get the correct result of an XQuery query. Our algorithm is called *BuildPlanCore*; it invokes *BuildPlanPred* to build sub-plans for query predicates.

5.1 Building FP Plan for Query Predicate

A function *BuildPlanPred* building an FP plan for a query predicate is shown in Algorithm 2. The algorithm is called with a parameter q_o – an output predicate node. In Line 1, a plan P_{q_o} is set to an IS operator. Lines 2 – 16 iterate through all neighbors[1] q of q_o and recursively build a plan P_q (Line 3) for q without the neighbor q_o. Then a binary semi-join operator is used to restrict the result of the plan P_{q_o}. The result is restricted such that it has to have an ancestor (Line 6), a parent (Lines 8), a descendant (Line 12) or a child (Line 14) in the result of P_q. Let us note that the function `rel(a, b)` returns the XPath relationship between a and b.

[1]By neighbors of a query node q we mean all its children and its parent (if q is not the root).

Algorithm 2: BuildPlanPred(q_o)

```
   input  : Query node q_o
   output: FP plan P_{q_o} for q_o
1  P_{q_o} ← IS(q_o);
2  foreach q ∈ neighbors of q_o do
3      P_q ← BuildPlanPred(q without q_o);
4      if q = parent(q_o) then
5          if rel(q_o, q) = AD then
6              P_{q_o} ← FilterDescAD(P_q, P_{q_o});
7          else if rel(q_o, q) = PC then
8              P_{q_o} ← FilterDescPC(P_q, P_{q_o});
9
10     else if q_o = parent(q) then
11         if rel(q_o, q) = AD then
12             P_{q_o} ← FilterAncAD(P_{q_o}, P_q);
13         else if rel(q_o, q) = PC then
14             P_{q_o} ← FilterAncPC(P_{q_o}, P_q);
15
16
17 return P_{q_o};
```

In Line 2, we iterate through the neighbors in an arbitrary order. We can, for example, use the order defined by the query syntax. However, it is clear that another order produces another plan. In [26], they try all permutations and select the cheapest one, i.e., a cost-based optimization is used. In our experiment described in Section 7.4, we show that the performance of plans for all permutations do not vary too much. In other words, the performance of any plan is close to the performance of the optimal plan, which is caused by utilizing only the semi-joins instead of partial-joins.

EXAMPLE 5.1. *Let us explain how the algorithm BuildPlanPred builds a plan for a sample predicate* $\Pi_{\#d}$ *in Figure 6a. The algorithm starts with the output predicate node* #d *and it first recursively builds plans for* #a *and* #e*, without* #d *itself (see Figure 6b and Figure 6c).* #a *and* #e *in the recursive call are treated as output predicate nodes.*

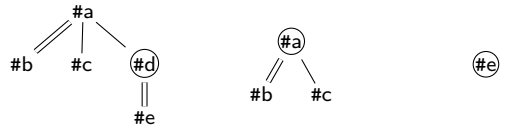

(a) Predicate $\Pi_{\#d}$ (b) Subtree in a recursive call for #a (c) Subtree in a recursive call for #e

Figure 6: Example of a predicate

The plan of #a *(without* #d *as a child) is computed such that we first recursively build plans for* #b *and* #c *without* #a *as a neighbor. The plans for* #b *and* #c *are* IS(#b) *and* IS(#c)*, respectively. The plan for* #a *is then:*

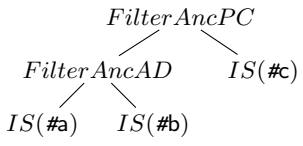

The plan of #e *(without* #d *as the parent) is* $IS($#d$)$*. Finally, the plan of the predicate* $\Pi_{\#d}$ *is as follows:*

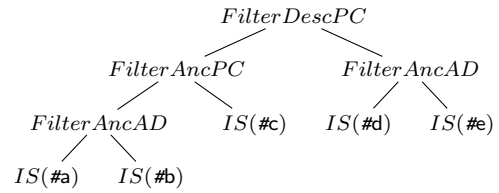

5.2 Building FP Plan for Query Core

A function *BuildPlanCore* building an FP plan for a query core is shown in Algorithm 3. The output is a plan producing a table where each column corresponds to one output query node. The order of columns corresponds to the pre-order traversal of the output query nodes. To meet the XQuery semantics, the table is lexicographically sorted by the columns.

Algorithm 3: BuildPlanCore(q)

```
   input : Query node q
   output: FP plan P for q
1  P ← BuildPlanPred(Π_q);
2  foreach c ∈ core_children(q) do
3      α ← rel(q, c);
4      if isOutput(c) then
5          R ← BuildPlanCore(c);
6          P ← StackTreeAnc_['1*','1*',1,1,α](P, R);
7      else
8          R ← BuildPlanPred(Π_c);
9          i ← 1;
10         d ← c;
11         while ¬ isOutput(d) do
12             d ← core_child(d);
13             if ¬ isOutput(d) then
14                 S ← BuildPlanPred(Π_d);
15             else
16                 S ← BuildPlanCore(d);
17             β ← rel(d, parent(d));
18             R ← StackTreeDesc_['10*','1*',i,1,β](R, S);
19             i ← 2;
20         P ← StackTreeAncSrt_['1*','01*',1,2,α,1](P, R);
21 return P;
```

The algorithm is invoked with the root query node of the query core as a parameter. In Line 1, it builds a plan P for a predicate of the root query node of the query core. In Lines 2 – 20, the algorithm loops through all child core query nodes c of q and successively extends the plan P. The function `core_children(`q`)` returns all core query nodes that are children of q.

If c is an output query node, a plan R for c is recursively built (Line 5) and it is joined with P using *StackTreeAnc* (Line 6). A more complicated situation becomes if c is a non-output query node. A plan R for the predicate of c is built (Line 8) and the child core query node of c is assigned to d (Line 12). The function `core_child(`q`)` returns the first child core query node of q. It is guaranteed that if q is a non-output core query node, it has exactly one child core query node (we discuss this property in Appendix B). Now, if d is still a non-output query node, a plan S is built using *BuildPlanPred* (Line 14). Otherwise, *BuildPlanCore* is recursively invoked (Line 16). In Line 18, the plans R and S are joined using *StackTreeDesc*. According to the variable i, the first or the second column of R is used as the join column.

We repeat the navigation to the non-output core query nodes (Lines 11 – 19) until an output query node is found. In Line 20, R and P are joined using *StackTreeAncSrt*. The secondary relationship test (see Section 4.1) with the first column of R ($k = 1$) ensures that the output satisfies the required relationship α.

Unlike from *BuildPlanPred*, *BuildPlanCore* has to iterate through child query nodes (Lines 2 – 20) in the order defined by the query syntax, so the result is sorted by output query nodes in a pre-order traversal (see Section 2.2). In [26], they compute the cost of all orders and pick the cheapest one which, however, does not have to produce the sorted result. Therefore, our algorithm is deterministic and, for a query core, it generates exactly one FP plan; it is not necessary to utilize any cost-based optimization technique.

Example 5.2. *Let us demonstrate how the algorithm BuildPlanCore builds a plan for the query core in Figure 2a. The algorithm starts with the parameter* $q =$ #a*, since* #a *is the root of the query core. The first step is to compute a plan* $P_{\#a}$ *for the predicate* $\Pi_{\#a}$ *(Line 1).*

After that we loop through the children of #a*.* #b *is the first one of them and it is a non-output query node. This is the more complicated case handled by Lines 8 – 20 in Algorithm 3. We start with building of a plan* R *for the predicate of* #b *(Line 8). Since the predicate of* #b *is the query node itself, the plan* $R = IS($#b$)$*. In Line 12,* d *is set to* #c*. Since* #c *is an output query node, we recursively invoke BuildPlanCore to compute a plan* S*. Since* #c *does not have any children, the plan* $S = IS($#c$)$*. In Line 18, we join the plans* R *and* S *using StackTreeDesc, the plan* R *is now:*

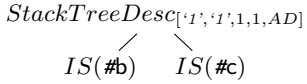

Such a plan produces a table of two columns for the query nodes #b *and* #c *ordered by the second column, respectively. Since the variable* d *is now an output query node, the loop in Line 11 is terminated and we move to Line 20. We join* P *with* R *using StackTreeAncSrt:*

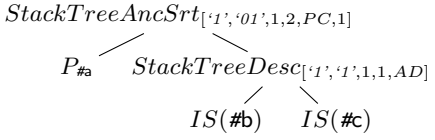

The mask M_d *is set to '01' since we want to retain only the second column corresponding to* #c *computed by* R*, because* #b *is a non-output query node. The indexes* i *and* j *are set to 1 and 2, respectively, since we need to join data nodes of* #a *with the data nodes of* #c*. We use the secondary relationship test to ensure that the data nodes of* #b *are in the PC relationship with* #a*.*

Finally, the plan P *is extended by the plan* $P_{\#d}$ *of the predicate for* #d*. The result of the algorithm is the following plan:*

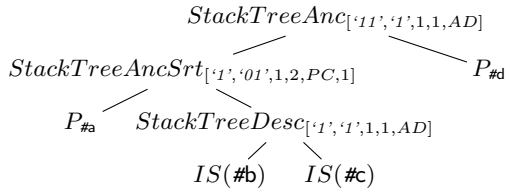

6. ANALYSIS OF THE BINARY JOIN APPROACH

For the purposes of our analysis, let us define the following properties of a stream. A *recursive stream* is a stream with some data nodes having the AD relationship. Recursive streams are caused by recursive tags in an XML document. A *repeating stream* is a stream with repeating data nodes, i.e., a repeating stream can return a particular data node multiple times. A *simple stream* is a non-recursive and non-repeating stream.

DEFINITION 1. *(Optimal TPQ processing approach) If an approach is optimal for a TPQ, then the time complexity is linear with respect to the size of the input and output and the space complexity is linear with respect to d, where d is the depth of an XML document.*

Until now the optimality has been proven only for the holistic approaches. Their optimality is proven with respect to the TPQ properties [7, 9] (e.g., all holistic approaches are optimal for a TPQ having only AD relationships) or with respect to the recursiveness of input streams [3].

DEFINITION 2. *(The linear time complexity condition) A binary join algorithm satisfies the linear time complexity condition if $|T_o| \leq |T_a|+|T_d|$, where $|T_o|$ is the size of the output and $|T_a|$ and $|T_d|$ are sizes of the inputs.*

LEMMA 6.1. *(The linear time complexity of the binary join approach for a TPQ Q). Let $\sum |T_i|$ be the sum of sizes of all input streams defined by query nodes in Q. The binary join approach to process Q has the linear time complexity with respect to $\sum |T_i|$ if each binary join in the corresponding FP plan satisfies the linear time complexity condition.*

PROOF. If each binary join in an FP plan satisfies the linear time complexity condition, then, in the worst case, for each binary join in an FP plan $|T_o| = |T_a|+|T_d|$. In such a case, for the root binary join $|T_a|+|T_d| = \sum |T_i|$. For each other binary join algorithm in the plan $|T_a|+|T_d| < \sum |T_i|$. Therefore, since (1) for each binary join algorithm in an FP plan $|T_a| + |T_d| \leq \sum |T_i| \wedge |T_o| \leq \sum |T_i|$ and (2) a binary join algorithm works with the linear time complexity with respect to the size of its input and output [1], the binary join approach has the linear time complexity with respect to $\sum |T_i|$. □

The linear time complexity condition is naturally satisfied for all semi-joins since for *FilterAncAD* and *FilterAncPC* $|T_o| \leq |T_a|$ and for *FilterDescAD* and *FilterDescPC* $|T_o| \leq |T_d|$. It is also satisfied for the partial-joins *StackTreeDesc* and *StackTreeAnc* when T_a is simple since each data node from T_d is joined with at most one data node from T_a.

DEFINITION 3. *(The linear space complexity condition) A binary join algorithm satisfies the linear space complexity condition if it has the linear space complexity with respect to d.*

LEMMA 6.2. *(The space complexity of the binary join approach for a TPQ Q). The binary join approach has the linear space complexity with respect to d if each binary join algorithm in an FP plan corresponding to Q satisfies the linear space complexity condition.*

PROOF. If each binary join algorithm in an FP plan satisfies the linear space complexity condition, then each binary join works with the linear space complexity with respect to d. Since d is the same for each algorithm in an FP plan, the binary join approach works with the linear space complexity with respect to d. □

Let us discuss which binary join algorithms satisfy the linear space complexity condition. We start with the semi-joins *FilterAncAD* and *FilterDescAD*. Since these semi-joins do not use any stack or lists (see Algorithm 1) they have the constant space complexity. The space complexity of other binary joins is affected mainly by the type of the stream T_a.

- If T_a is simple, there is at most one entry in the stack of partial-joins *StackTreeDesc* and *StackTreeAnc*. In such a case, for *StackTreeAnc* it is also not necessary to use the self- or inherited-lists [1]. Therefore, in this case, the space complexity of the partial-joins is constant. The same property is valid for semi-joins *FilterAncPC* and *FilterDescPC* since they are derived from the partial-joins.
- If T_a is recursive, the partial-join *StackTreeDesc* and the semi-join *FilterDescPC* can store more data nodes from T_a on the stack, however there are no self- or inherited-lists and the size of the stack is bounded by d. The partial-join *StackTreeAnc* and the semi-join *FilterAncPC* have to use the self- and inherited-lists and the size of these lists is not bounded by d. We discuss an example of such a situation in Appendix C.
- If T_a is repeating, the size of the stack of the partial-joins *StackTreeDesc* and *StackTreeAnc* is not bounded by d.

Table 3 summarizes which binary join algorithms satisfy the linear time complexity condition (TC) (Definition 2) and the linear space complexity condition (TS) (Definition 3) with respect to the properties of T_a. Using the algorithms *BuildPlanPred* and *BuildPlanCore*, all semi-joins in an FP plan are always processed before all partial-joins. Since only the partial-joins can produce repeating streams, we do not analyse properties of semi-joins if T_a is repeating. We also do not explicitly analyse properties of the novel *StackTreeAncSrt* (see Section 4.1) since it has the same properties as *StackTreeAnc*.

Algorithm		T_a **is**					
		simple		**recursive**		**repeating**	
		TC	SC	TC	SC	TC	SC
semi-joins	*FilterDescAD*	✓	✓	✓	✓	–	–
	FilterAncAD	✓	✓	✓	✓	–	–
	FilterDescPC	✓	✓	✓	✓	–	–
	FilterAncPC	✓	✓	✓	✗	–	–
part.-joins	*StackTreeDesc*	✓	✓	✗	✓	✗	✗
	StackTreeAnc	✓	✓	✗	✗	✗	✗

Table 3: Linear time and space complexity conditions of the binary join algorithms for T_a

Using the Table 3 and Lemmas 6.1 and 6.2, we can now derive several classes of queries for which the binary join approach is optimal.

LEMMA 6.3. *(Optimal predicate processing). The binary join approach is optimal to process a predicate if there is not an edge with the following properties: (i) having the PC relationship, and (ii) not lying on a path from the root to the output predicate node, and (iii) having the recursive stream of the parent query node.*

PROOF. A predicate is processed only by the semi-join algorithms (see Algorithm 2). An edge with the PC relationship not lying on a path between the root and the output predicate node leads to the usage of *FilterAncPC*. This is the only semi-join algorithm not having the linear space complexity with respect to d if T_a is recursive. □

THEOREM 1. *(TPQ with one output query node). The binary join approach is optimal to process a TPQ with one output query node if there is not an edge: (i) with the PC relationship, and (ii) not lying on a path from the root to the output query node, and (iii) having the recursive stream of the parent query node.*

PROOF. Theorem 1 directly follows Lemma 6.3 since each TPQ with one output query node is a predicate itself. □

THEOREM 2. *(TPQ with more output query nodes). The binary join approach is optimal for a TPQ with more output query nodes if (i) it is optimal for all predicates, and (ii) the query core is a path, and (iii) the streams of the core query nodes are not recursive.*

PROOF. Let us first prove Theorem 2 for a query where all core query nodes are output. Each edge in the query core is then processed by *StackTreeAnc* (see Algorithm 3). If the query core is a path, T_a for each *StackTreeAnc* is computed by semi-joins of a predicate, therefore T_a is not repeating. Only T_d can be computed by another *StackTreeAnc* and, therefore, it can be repeating. Since T_a is not repeating and not recursive, it is simple.

If there are non-output query nodes, each edge in the query core, where the parent is a non-output query node, is processed by *StackTreeDesc* (see Algorithm 3). The output stream T_o of *StackTreeDesc*, where T_a and T_d are simple, is also simple since a data node from T_d can be joined with at most one data node from T_a. Since T_a of *StackTreeDesc* is the output of (a predicate or another *StackTreeDesc*), T_a is simple. If T_a is simple for each partial-join in an FP plan, the binary join approach is optimal. □

EXAMPLE 6.1. *In Figure 7a, there is a query Q4 without non-output query nodes. The whole query is the query core itself and it is a path. The corresponding plan is depicted in Figure 7b. Let us consider that the streams of all query nodes are non-recursive. In such a case, the inputs of all StackTreeAnc partial-joins are simple. Therefore, the binary approach processing the FP plan for Q4 is optimal.*

EXAMPLE 6.2. *The query Q5 in Figure 8a is also a query core itself and it is a path, however,* #b *and* #c *are non-output query nodes. The corresponding FP plan is depicted in Figure 8b. T_a of $StackTreeDesc_2$ is the output of $StackTreeDesc_1$. Assuming that the streams for* #b *and* #c *are non-recursive, each data node of* #c *is joined with at most one data node of* #b*. Therefore the output of $StackTreeDesc_1$ is non-recursive and non-repeating, i.e., it is simple. Since T_a of all partial-joins are simple, the binary approach processing the FP plan for Q5 is optimal.*

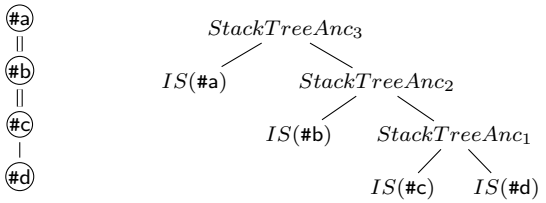

(a) Query **Q4** (b) FP plan of Q4

Figure 7: A query core without non-output query nodes

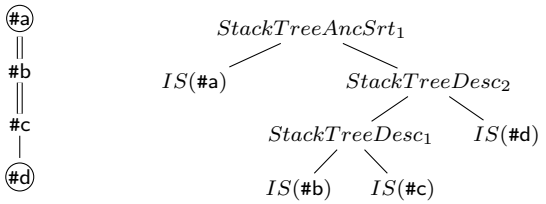

(a) Query **Q5** (b) FP plan of Q5

Figure 8: A query core with non-output query nodes

Let us compare the optimality of our binary join approach and the holistic join approaches. Every holistic join is optimal when a query contains only AD relationships [7]. A holistic join approach is optimal also for a TPQ if all streams corresponding to the query nodes having the PC relationship are non-recursive [3]. When compared to the novel optimality of binary join approach, *the only difference is that the optimality of holistic approaches is not influenced by the number of non-output query nodes.*

7. EXPERIMENTAL RESULTS

In this section, we experimentally demonstrate weak and strong aspects of binary and holistic join approaches under certain conditions. Both holistic join (HJ) and binary join (BJ) approaches are implemented in our native database management system called RadegastXDB (RXDB)[2]. Both approaches use the same index data structure (see Section 2.5) and no cost-based optimizer is included. Let us recall that we select the *GTPStack* holistic join [3] as a representative of the state-of-the-art holistic algorithms.

We employed three typical data collections with various characteristics and sizes, these collections are commonly used in many approaches [17, 10, 12, 3]. XMark [19] is a well-known XML benchmark representing a synthetic database of auctions. An XML document of auctions can be generated with the factor f that controls the size of the document; we used two factors: $f = 1$ and $f = 10$. The structure of the document is rather regular with a few recursive tags. SwissProt [22] is a real database of protein sequences with annotations; its structure is quite irregular without any recursive tags. TreeBank [22] is a real database of partially

[2] `http://db.cs.vsb.cz/Projects`

encrypted English sentences; it is characterized by a highly recursive and irregular structure. TreeBank is the smallest collection, but it is the most complicated collection from the query processing perspective due to its highly recursive nature. An overview of the most important statistics of these collections can be found in Table 4. We also show the size of the index data structures in RXDB.

Collection	Size (MB)	Number of data nodes	Max depth	Index size (MB)
XMark (f=1)	111	2,048,193	14	31
XMark (f=10)	1,137	20,532,805	14	315
SwissProt	109	5,166,890	7	79
TreeBank	82	2,437,667	38	37

Table 4: Statistics of data collections

For each data collection, we generated 50 unique *twig patterns* combining the AD and PC relationships. By a twig pattern we mean a TPQ without a specification of non-output and output query nodes. Some important statistics of these twig patterns are shown in Table 5, the complete list of the twig patterns can be found in Appendix E. For each twig pattern p, we randomly generated n queries where n is the number of query nodes in p, such that each query has a different number of output query nodes. Finally, there are 314 queries for XMark, 291 queries for SwissProt, and 351 queries for TreeBank. The twig patterns were generated such that each setting of output or non-output query nodes produces a non-empty result, i.e., the *selectivity* is greater than zero. We compute the selectivity of a query by the following formula:

$$\sigma = \frac{N_{out}}{N_{in}},$$

where N_{out} is the total number of distinct data nodes in the result of the query and N_{in} is the total number of data nodes in streams specified by output query nodes.

Statistic	Value
Min. query nodes	3
Max. query nodes	12
Min. tree depth	3
Max. tree depth	9
Max. children per query node	3

Table 5: Statistics of the tested queries

The experiments were performed on a machine with Intel Xeon E5-2690@2.9 GHz processor and the Microsoft Windows Server 2008 R2 Datacenter (SP1) operating system. During the experiments, the index data structures of RXDB were completely loaded in the main memory.

7.1 Non-output and Output Query Nodes

The purpose of this experiment is to show how the number of non-output and output query nodes affects the processing time of HJ and BJ. A typical example is demonstrated in Figure 9. For the demonstration, we used the twig pattern TB18 with 7 query nodes, therefore, 7 queries with the different number of output query nodes have been tested.

We can see that the processing time of HJ is nearly the constant, while the processing time of BJ depends on the number of output query nodes. For a query with one output query node, BJ uses only semi-join algorithms and it

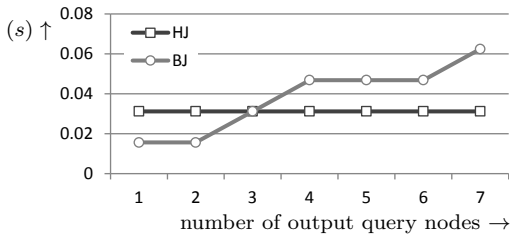

Figure 9: Comparison of HJ and BJ on a sample query (TB18)

outperforms HJ. On the other hand, when all query nodes are set to output, only partial-joins are used and HJ outperforms BJ.

Summarized results for all collections can be seen in Figure 10. Horizontal axis represents the ratio $\rho = \frac{n_o}{n}$, where n_o is the number of output query nodes and n is the total number of query nodes in a query. The values on the vertical axis represent the relative speed-up of BJ compared to HJ; it is computed as $\frac{t_{HJ}}{t_{BJ}}$, where t_{HJ} and t_{BJ} are total processing times of queries with ρ in the specific range using HJ and BJ, respectively. BJ outperforms HJ when $\frac{t_{HJ}}{t_{BJ}} > 1$. We can see that the lower value of ρ means the lower processing time of BJ compared to HJ. This is consistent with the analysis in Section 6: the more output query nodes there are in a query, the more repeating streams appear in the corresponding FP plan, and therefore, the binary join approach becomes non-optimal. On the other hand, the optimality of HJ does not depend on the number of output query nodes.

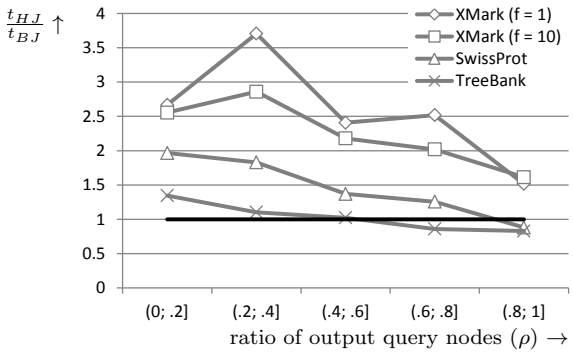

Figure 10: Comparison of HJ and BJ on queries with various ratios of output query nodes

The Figure 10 also compares our binary join approach with another binary join aproach presented in [26], where all query nodes are always output. The relative speed up of [26] to HJ can be therefore approximated by $\frac{t_{HJ}}{t_{BJ}}$ where $\rho \in (.8; 1]$. For such a range of ρ, the performance of BJ is the worst. We can see that considering the non-output query nodes means a more efficient query processing.

7.2 Intermediate Result Sizes

In this experiment, we focus on the size of the intermediate storage of the binary join approach, since it is often identified as a weak aspect of binary join approaches in the literature. By the intermediate storage we mean all dynamic data structures used by binary joins in an FP plan: stacks,

inherited-lists, and self-lists. For each query, we measured μ which represents the maximum number of data nodes in the intermediate storage.

Summarized results are shown in Figure 11. For five ranges of ρ in the horizontal axis, we computed arithmetic means of μ for each data collection separately. We can observe that for a small $\rho \in (0; .2]$, in the worst case, there are approximately 10 data nodes in the intermediate storage at one time of the query processing. For $\rho \in (.8; 1]$, the value of $\overline{\mu}$ is approximately 100 – 5,000. The higher ratio of output query nodes ρ means using of partial-joins in the FP plan, which causes repeating streams that have a negative impact on the space complexity (see Section 6).

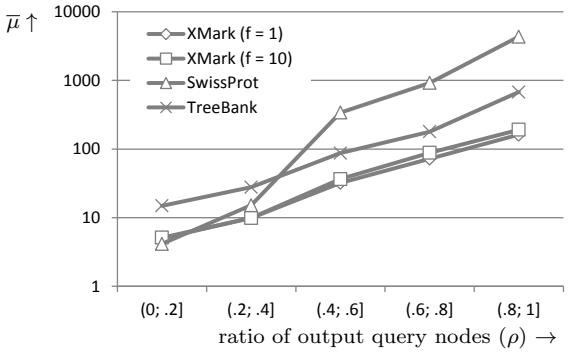

Figure 11: Maximum intermediate storage size for queries with variable ratio of output query nodes

In Table 6, we show $\max(\mu)$ of queries with one and all output query nodes. For the queries where $\mu = \max(\mu)$, we include the result size and the number of output query nodes n_o. In the last column, we show a ratio of $\max(\mu)$ to $n_o \times$ [result size], where $n_o \times$ [result size] means the total number of data nodes in the result[3]. The last column compares the intermediate result size to the final result size.

Collection	$\max(\mu)$	Result size	n_o	$\frac{\max(\mu)}{n_o \times [\text{result size}]}$ (%)
one output query node				
XMark (f=1)	13	273	1	4.76
XMark (f=10)	14	2,979	1	0.47
SwissProt	7	12,484	1	0.06
TreeBank	59	20	1	295.00
all output query nodes				
XMark (f=1)	902	111,676	6	0.13
XMark (f=10)	1,280	1,131,354	6	0.02
SwissProt	79,341	149,691	7	7.57
TreeBank	11,696	307	10	380.98

Table 6: Maximum intermediate storage sizes

For the queries with one output query node, we can see that the intermediate storage includes rather small numbers of data nodes. In the worst case, $\max(\mu) = 59$ (for TreeBank), it is approximately $3\times$ more than the number of data nodes in the result. Let us note that the number of data nodes in the result does not represent the byte size of the result. In this paper, the node label consists of 3 integer numbers, i.e., 12 bytes. In such a case, the byte size of the intermediate storage is never higher than $59 \times 12 = 708$ bytes.

For the queries where all query nodes are output, $\max(\mu)$ is significantly higher than in the case of queries with one output query node. In the worst case, $\max(\mu) = 79,341$ (approximately 1 MB of memory) for SwissProt, it is almost $4\times$ more than the number of data nodes in the result. Such a value shows a typical disadvantage of a binary join approach for some data collections and queries: the large intermediate result can be produced even if the final result is small. However, we can see that this behavior is observed only for one data collection in the experiment. Moreover, considering 12 bytes as the label size, then the worst case memory size necessary during the query processing is less then 1 MB which can be easily accommodated in the L2 cache of CPU.

The results also show that considering the non-output query nodes significantly reduces the intermediate storage size.

7.3 Query Selectivity

This experiment is focused on the processing time of BJ and HJ on queries with different selectivity. The results are presented in Figure 12, where the horizontal axis represents specific ranges of the selectivity and the values on the vertical axis represent the relative speed-up of BJ compared to HJ; it is computed as $\frac{t_{HJ}}{t_{BJ}}$.

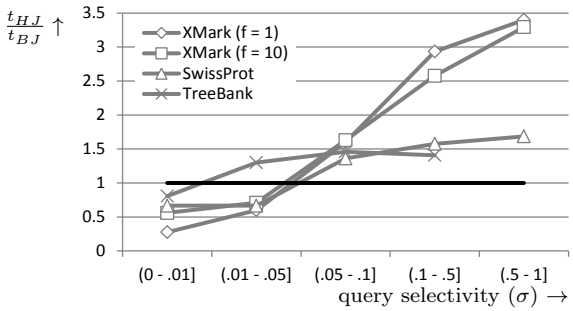

Figure 12: Comparison of HJ and BJ for queries with different selectivity

As a result, the lower the selectivity σ is, the more HJ outperforms BJ. On all data collections, HJ performs better when $\sigma < 0.01$. On the other hand, for higher values of σ, BJ outperforms HJ. For example, on the XMark collection (both factors), for queries where σ is close 1, BJ is almost $3.5\times$ faster than HJ. HJ performs better on high-selective queries, since it uses a more aggressive forwarding of the input streams of data nodes. For low-selective queries, BJ gives better results since it utilizes more simple logic to advance the input streams.

7.4 Query Plans

In Section 5, we proposed the algorithm *BuildPlanPred* to build an FP plan for a query predicate. The algorithm can iterate through children of a query node in an arbitrary order and, therefore, it can generate a number of FP plans. This experiment shows that the processing times of different FP plans for a query do not vary too much. We tested

[3]The result size of a query is the total number of output matches (tuples). To retrieve the total number of data nodes in the result, the result size must be multiplied by the number of output query nodes.

queries with one output node for TreeBank where at least 12 different FP plans could be generated.

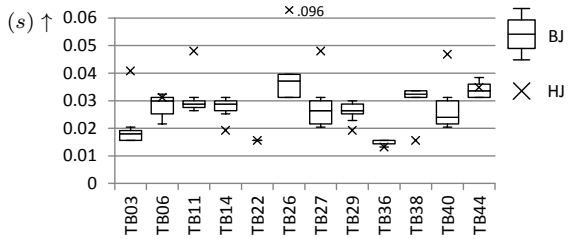

Figure 13: Processing different query plans

The results are depicted in Figure 13 with box plots of the processing times of BJ using different FP plans (the processing times of HJ are also included). We can see that the processing times of different plans vary within a few (less than 15) milliseconds. As we can observe in most of the queries, the whole box plots are either above or bellow the time of HJ, which means that the processing time of BJ is usually either better or worse than the processing time of HJ, no matter what plan is used. We repeated the same experiment over the other collections (XMark and SwissProt) with similar results. The results show that it is possible to use *BuildPlanPred* even without any cost-based optimizer.

7.5 Comparison of Native XML DBMSs

Finally, we compared our native XML database management system RXDB with other commonly used systems: BaseX[4] (BX) and MonetDB with an XQuery front-end[5] (MDB). In the experiment, we compared our novel algortihms with other existing implementations of a TPQ processing. Summarized results are shown in Table 7: the values represent total processing times of all testing queries in seconds. We can see that the best overall performance on all data collections is given by our RXDB using the binary join approach (R-BJ); R-BJ is approximately $2-5\times$ faster than MDB and by an order of magnitude faster than BX. R-HJ gives a worse overall performance compared to R-BJ, but it can process particular queries with a high ratio of output query nodes or highly-selective queries more efficiently (see Section 7.1 and Section 7.3, respectively).

Collection	BX	MDB	R-BJ	R-HJ
XMark (f=1)	27.7	15.2	**3.2**	6.3
XMark (f=10)	280.1	62.0	**33.3**	64.6
SwissProt	206.6	57.4	**12.7**	15.9
TreeBank	129.7	29.7	**11.2**	11.6

Table 7: Total query processing times (s)

Table 8 shows processing times of queries processed on XMark ($f = 10$) and TreeBank listed in Appendix F. The first 5 queries on both collections have one output query node, the other 5 queries have all output query nodes. The lowest time for each query is highlighted. We can see that the relative speed-up of R-BJ and R-HJ over MDB or BX is similar for all the queries. However, for queries with all output nodes on TreeBank, R-HJ gives lower processing times than R-BJ, which supports our analysis in Section 6.

[4] `http://basex.org/`

[5] `https://www.monetdb.org/XQuery/`

	XMark (f=10)				TreeBank			
	BX	MDB	R-BJ	R-HJ	BX	MDB	R-BJ	R-HJ
1	**.023**	.037	**.023**	**.023**	.285	.062	**.012**	.016
2	.284	.131	**.031**	.109	.499	.109	**.023**	.043
3	**.030**	.059	.060	.057	.093	.051	**.008**	.014
4	.736	.351	**.062**	.224	.306	.109	**.031**	**.031**
5	.376	.183	**.047**	.220	.251	.069	**.031**	.057
6	1.218	.351	**.255**	.310	.622	.118	**.094**	.098
7	1.345	**.126**	.127	.199	.618	.110	.031	**.021**
8	.874	.167	**.092**	.197	.397	.071	.021	**.020**
9	.463	.059	**.047**	.107	1.117	.281	.045	**.027**
10	1.263	.346	**.242**	.472	.567	.101	.059	**.047**

Table 8: Processing times of particular queries (s)

8. CONCLUSION

In this article, we introduced a new binary join approach processing a TPQ with defined output and non-output query nodes using an FP plan; considering output and non-output query nodes supports the XQuery semantics, which leads to a more efficient query processing in many cases. We also showed that it is not necessary to integrate a cost-based optimizer in the approach since the performance of all individual plans is rather similar. Moreover, we proved that for a certain class of TPQs such an FP plan has the linear time complexity with respect to the size of the input and output and the linear space complexity with respect to the XML document depth (i.e., the same complexity as the holistic join approaches). We experimentally compared the binary approach with a state-of-the-art holistic approach (*GTPStack*) for 4 data collections with different characteristics: although we showed that our binary approach overcomes *GTPStack* when the average processing time is considered, we put forward that *GTPStack* performs more efficiently for the low-selective queries for all data collections (up-to the selectivity 0.05) and for the queries with a higher ratio of non-output query nodes for two data collections. As a result, an optimal XQuery processor should include both types of approaches to select one of them considering the query selectivity and the ratio of non-output query nodes.

In the experiments, we also compared the binary approach with another FP plan-enabling binary approach and with other commonly used XML database management systems. In the case of the comparison with another binary approach, we mainly showed that the consideration of non-output query nodes increases the query processing time up-to $2.5\times$. When compared with the other XML database management systems, our approach lowers the average query processing time from $2\times$ to $17\times$. In the experiments, we also showed that the intermediate result sizes of our binary approach did not exceed 1 MB in the worst case although the high intermediate result sizes of binary join approaches are often criticized in the literature.

9. REFERENCES

[1] S. Al-Khalifa, H. Agadish, N. Koudas, J. M. Patel, D. Srivastava, and Y. Wu. Structural joins: A primitive for efficient xml query pattern matching. In *Data Engineering, 2002. Proceedings. 18th International Conference on*, pages 141–152. IEEE, 2002.

[2] S. Al-Khalifa and H. Jagadish. Multi-level operator combination in xml query processing. In *Proceedings of the eleventh international conference on*

Information and knowledge management, pages 134–141. ACM, 2002.

[3] R. Bača, M. Krátký, T. W. Ling, and J. Lu. Optimal and efficient generalized twig pattern processing: a combination of preorder and postorder filterings. *The VLDB Journal*, 22(3):369–393, 2013.

[4] R. Bača, P. Lukáš, and M. Krátký. Cost-based holistic twig joins. *Information Systems*, 52:21–33, 2015.

[5] M. Brantner, S. Helmer, C.-C. Kanne, and G. Moerkotte. Full-fledged Algebraic XPath Processing in Natix. In *Proc. of the 21st Int. Conf. on Data Engineering (ICDE)*, pages 705–716, 2005.

[6] M. Brantner, S. Helmer, C.-C. Kanne, and G. Moerkotte. Full-fledged algebraic XPath processing in Natix. In *Proceedings of Data Engineering, 2005. ICDE 2005.*, pages 705–716. IEEE, 2005.

[7] N. Bruno, N. Koudas, and D. Srivastava. Holistic twig joins: optimal xml pattern matching. In *Proceedings of the 2002 ACM SIGMOD international conference on Management of data*, pages 310–321. ACM, 2002.

[8] S. Chen, H.-G. Li, J. Tatemura, W.-P. Hsiung, D. Agrawal, and K. S. Candan. Twig2Stack: Bottom-up Processing of Generalized-tree-pattern Queries Over XML documents. In *Proceedings of VLDB 2006*, pages 283–294, 2006.

[9] T. Chen, J. Lu, and T. W. Ling. On Boosting Holism in XML Twig Pattern Matching Using Structural Indexing Techniques. In *Proceedings of ACM SIGMOD 2005*, pages 455–466. ACM Press, 2005.

[10] Z. Chen, H. V. Jagadish, L. V. S. Lakshmanan, and S. Paparizos. From tree patterns to generalized tree patterns: on efficient evaluation of XQuery. *Proceedings of the 29th international conference on Very Large Data Bases, VLDB 2003*, pages 237–248, 2003.

[11] H. Georgiadis, M. Charalambides, and V. Vassalos. Efficient physical operators for cost-based xpath execution. In *Proceedings of the 13th International Conference on Extending Database Technology*, pages 171–182. ACM, 2010.

[12] N. Grimsmo, T. A. Bjørklund, and M. L. Hetland. Fast Optimal Twig Joins. In *Proceedings of the 36th International Conference on Very Large Data Bases, VLDB 2010*, pages 894–905. VLDB Endowment, 2010.

[13] T. Grust, S. Sakr, and J. Teubner. XQuery on SQL Hosts. In *Proc. of the 29th Int. Conf. on Very Large Data Bases (VLDB)*, pages 252–263, 2004.

[14] T. Grust, M. van Keulen, and J. Teubner. Staircase Join: Teach a Relational DBMS to Watch Its (Axis) Steps. In *Proceedings of VLDB 2003*, pages 524–535, 2003.

[15] J. Lu, T. Chen, and T. W. Ling. Efficient Processing of XML Twig Patterns with Parent Child Edges: a Look-ahead Approach. In *Proceedings of ACM CIKM 2004*, pages 533–542. ACM Press, 2004.

[16] J. Lu, T. W. Ling, C.-Y. Chan, and T. Chen. From Region Encoding to Extended Dewey: on Efficient Processing of XML Twig Pattern Matching. In *Proceedings of the 31st International Conference on Very Large Data Bases, VLDB 2005*, pages 193–204, 2005.

[17] P. Michiels, G. Mihaila, and J. Siméon. Put a Tree Pattern in Your Algebra. In *Proceedings of the 23th International Conference on Data Engineering, ICDE 2007*, pages 246–255, 2007.

[18] S. Paparizos, Y. Wu, L. V. S. Lakshmanan, and H. V. Jagadish. Tree logical classes for efficient evaluation of XQuery. In *Proceedings of the 2004 ACM SIGMOD international conference on Management of data*, pages 71–82. ACM, 2004.

[19] A. Schmidt, F. Waas, M. Kersten, M. J. Carey, I. Manolescu, and R. Busse. Xmark: A benchmark for xml data management. In *Proceedings of the 28th international conference on Very Large Data Bases*, pages 974–985. VLDB Endowment, 2002.

[20] S. Son, H. Shin, and Z. Xu. Structural semi-join: A light-weight structural join operator for efficient xml path query pattern matching. In *Database Engineering and Applications Symposium, 2007. IDEAS 2007. 11th International*, pages 233–240. IEEE, 2007.

[21] I. Tatarinov, S. D. Viglas, K. Beyer, J. Shanmugasundaram, E. Shekita, and C. Zhang. Storing and querying ordered xml using a relational database system. In *Proceedings of the 2002 ACM SIGMOD international conference on Management of data*, pages 204–215. ACM, 2002.

[22] University of Washington Database Group. The XML Data Repository, `http://www.cs.washington.edu/research/xmldatasets/`, 2002.

[23] A. M. Weiner and T. Härder. Using structural joins and holistic twig joins for native xml query optimization. In *East European Conference on Advances in Databases and Information Systems*, pages 149–163. Springer, 2009.

[24] A. M. Weiner and T. Härder. An integrative approach to query optimization in native XML database management systems. In *Proceedings of the Fourteenth International Database Engineering & Applications Symposium*, IDEAS '10, pages 64–74, New York, NY, USA, 2010. ACM.

[25] H. Wu, T. W. Ling, and G. Dobbie. Tp+ output: modeling complex output information in xml twig pattern query. In *International XML Database Symposium*, pages 128–143. Springer, 2010.

[26] Y. Wu, J. M. Patel, and H. Jagadish. Structural join order selection for xml query optimization. In *Data Engineering, 2003. Proceedings. 19th International Conference on*, pages 443–454. IEEE, 2003.

[27] M. Yoshikawa, T. Amagasa, T. Shimura, and S. Uemura. XRel: a Path-based Approach to Storage and Retrieval of XML Documents Using Relational Databases. *ACM Transactions on Internet Technology*, pages 110–141, 2001.

[28] C. Zhang, J. Naughton, D. DeWitt, Q. Luo, and G. Lohman. On supporting containment queries in relational database management systems. In *ACM SIGMOD Record*, volume 30, pages 425–436. ACM, 2001.

APPENDIX

A. PROBLEMATIC ORDER OF FOR CLAUSES

In Figure 14a, we see a query Q6 for which there is not any TPQ whose pre-order traversal of the output query nodes matches the order of the 'for' clauses. If we processed Q6 by a TPQ in Figure 14b, the result would be sorted by (#a, #b, #d, #c) and an additional sorting would be necessary. To the best of our knowledge, in the literature, there is no binary or holistic join approach with focus on the processing of such queries without an additional sorting of the result.

```
for $a in //a
for $b in $a/b
for $c in $a/c
for $d in $b/d
return ($a, $b, $c, $d)
```

(a) Query **Q6**

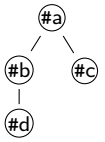

(b) TPQ of Q6

Figure 14: An XQuery query with a problematic order of 'for' clauses

B. PROPERTIES OF A QUERY CORE

In this section, we prove that in a query core, each non-output query node has exactly one child. This follows from a property of a real XQuery query: for each pair of output query nodes, their nearest common ancestor is also an output query node.

We give a proof by contradiction. Let q_i be a non-output core query node. Let q_a and q_b be core query nodes such that q_i is their parent. Then q_a and q_b lie on a path between two output query nodes (see the definition of a query core in Section 2). Therefore, in subtrees of q_a and q_b there must be output query nodes q_{o_a} and q_{o_b}, respectively. However, q_i is the nearest common ancestor of q_{o_a} and q_{o_b} and, therefore, it has to be an output query node.

C. SUBOPTIMAL QUERY PROCESSING FOR RECURSIVE STREAM

In this section, we discuss an example where the binary join approach is not optimal to process a query because of a recursive stream T_a and the unbounded size of the self- and inherited-lists of the *StackTreeAnc* algorithm. Let us process the query Q7 in Figure 15b on the document in Figure 15a. The plan of Q7 is depicted in Figure 15c and we see that it consists only of one partial-join algorithm *StackTreeAnc*.

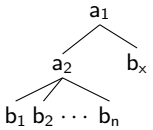

(a) XML document

(b) Query **Q7**

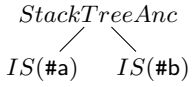

(c) Plan of Q7

Figure 15: Suboptimal query processing

The operator starts with reading the data nodes a_1 and a_2. Since a_2 is a descendant of a_1, the stream T_a is recursive and both data nodes are put on the stack. After that, the algorithm reads the data nodes $b_1 \cdots b_n$ and it determines that all pairs $[a_2, b_1] \cdots [a_2, b_n]$ should be output. However, these pairs cannot be output until a_1 is popped out from the stack. This is because of the possibility that there is a data node b_x following a_2 that will be joined with a_1. Since *StackTreeAnc* is an ancestor-sorted partial-join, it has to produce the result sorted by ancestors, i.e., it has to first output all pairs joined with a_1. Therefore, $[a_2, b_1] \cdots [a_2, b_n]$ are not output immediately. Instead, they are placed into the self-list of a_2. The situation is depicted in Figure 16.

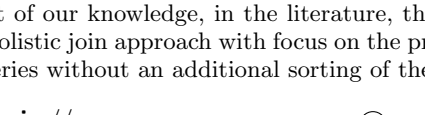

Figure 16: StackTreeAnc after reading of b_n

After the reading of b_n, a_2 can be popped out from the stack and its self-list is added into the inherited-list of a_1. Subsequently, b_x is read and $[a_1, b_x]$ is then added into the self-list of a_1. Finally, a_1 is popped out from the stack and its self-list and inherited-list are enumerated to the output, respectively.

In Figure 16, we can see that there are possibly many pairs $[a_2, b_1] \cdots [a_2, b_n]$ that have to be stored in the self-list of a_2. The number of pairs is not bounded by the depth of the document, therefore, the binary join approach is not optimal to process the query Q7 on the document in Figure 15a.

D. PROCESSING OF GTP

In [10] a *generalized twig pattern* (GTP) is introduced. Such a query model supports the XQuery semantics in many aspects. Apart from a specification of output and non-output query nodes, it defines aggregation query nodes which correspond to 'let' clauses. Moreover, it supports boolean formulas with 'or' and 'not' connectives. In this section, we discuss how to extend our binary join approach to support such a model. Let us note that RXDB (our prototype of a native XML database) supports the GTP.

D.1 'let' Clauses

Let us call *aggregation query nodes* the query nodes which correspond to the 'let' clauses. In Figure 17a, we see an example of such a query. The corresponding GTP is depicted in Figure 17b, where the aggregation query nodes are squared, i.e., #e is the only aggregation query node.

```
for $i in //a
for $j in $a/b//c
let $k := $a//d/e
return ($i, $j, $k)
```

(a) Query **Q8**

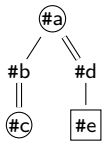

(b) GTP of Q8

Figure 17: An XQuery query with a 'let' clause

The processing of the aggregation query nodes is similar to the processing of output query nodes. However, in a principle, there are two differences:

1. In Section 2.4, we defined a table as an array of tuples of data nodes. When the aggregation query nodes are considered, the columns corresponding to such query nodes contain sets of data nodes. The result of Q8 on the document in Figure 1a is depicted in Table 9. We

see, the cells in the column #e contain sets of data nodes.

#a	#c	#e
a_1	c_1	e_1, e_2
a_1	c_2	e_1, e_2
a_2	c_3	e_3

Table 9: The result of the query Q8

2. Since the sets of data nodes corresponding to the aggregation query nodes can be empty, we need to introduce an outer partial-join. Otherwise we could loose some of the results of the 'for' clauses. In Line 20 of *BuildPlanCore* (see Algorithm 3), if c is an aggregation query node, we need to introduce an outer version of the ancestor-sorted partial-join *StackTreeAncSrt*. The implementation the partial-join is straightforward: all tuples of T_a are always retained, even if there are no corresponding tuples in T_d.

D.2 Boolean Formulas

The TPQ model defined in Section 2.1 can be used to process only XQuery queries with conjunctive boolean formulas, i.e., boolean formulas where only 'and' connectives appear. However, in many real-world queries, we can also find 'or' and 'not' connectives. An example of such a query can be found in Figure 18a. In the corresponding GTP (see Figure 18b), we see that #d is labeled with a boolean formula that has to be satisfied between #f and #c.

```
for $i in //a
for $j in $a//d[.//f or not(.//c)]
return ($i, $j)
```

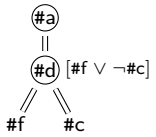

(a) Query **Q9** (b) GTP of Q9

Figure 18: An XQuery query with a 'let' clause

To process a query node q labeled with a boolean formula, we introduce a generalized version of *FilterAncPC* such that it can have more input streams T_d. Each T_d corresponds to one child query node of q, nevertheless, it still has one input stream T_a. Before a data node a of T_a is output, we check whether the boolean formula is satisfied between the corresponding data nodes of the streams T_d. We can observe that such an approach brings some similarities with a holistic join algorithm, since the generalized version of *FilterAncPC* processes more than two streams at once.

E. TESTING TWIG PATTERNS

For each twig pattern p in this appendix, we randomly generated n queries, where n is the number of query nodes in p, such that each query has a different number of output query nodes.

E.1 XMark

XM01 //categories/category[//description//text and /@id]
XM02 //closed_auctions/closed_auction[//price]/buyer/@person
XM03 //regions/samerica/item[/mailbox/mail/to]/incategory/@category
XM04 //open_auctions//open_auction[/reserve and //bidder/personref/@person]/itemref/@item
XM05 //open_auctions/open_auction[/interval//start]
XM06 //regions/namerica/item[//from]/incategory/@category
XM07 //regions[//europe]/namerica/item/description[/parlist]//listitem/text
XM08 //people/person[/emailaddress and //creditcard]/@id
XM09 //closed_auctions/closed_auction[/type]/quantity
XM10 //people/person/profile[//business]/interest/@category
XM11 //closed_auctions//annotation[//happiness and /description/text]
XM12 //closed_auctions//closed_auction[//date]//itemref[/@item]
XM13 //closed_auctions/closed_auction[//buyer]/itemref[/@item]
XM14 //closed_auctions/closed_auction[//seller/@person]/annotation
XM15 //open_auctions/open_auction[/annotation/description//text]//date
XM16 //open_auctions/open_auction//bidder[/increase]/date
XM17 //open_auctions/open_auction[/@id]/bidder/date
XM18 //regions//item[//text/keyword]/mailbox/mail[/to]/from
XM19 //regions[/samerica]//item[/description/text]/incategory[/@category]
XM20 //regions//item[//bold and /payment]//shipping
XM21 //regions//item[/description/text and /location]/shipping
XM22 //open_auctions//open_auction[/current and //bidder]/initial
XM23 //open_auctions/open_auction[/bidder]//interval
XM24 //open_auctions/open_auction[/bidder]//initial
XM25 //open_auctions//open_auction[//bidder[/time]/date]/initial
XM26 //open_auctions/open_auction[/bidder[/personref/@person]/time]/interval//end
XM27 //categories/category[/description/text//emph and /@id]
XM28 //categories/category[//name and /@id]/description//text
XM29 //categories/category[/name]/description/parlist/listitem/text
XM30 //people/person[/@id and /homepage and //name]
XM31 //open_auctions/open_auction[/reserve]//seller/@person
XM32 //categories/category[/description/text/keyword]/name
XM33 //regions//samerica/item[/incategory/@category and //name]//shipping
XM34 //open_auctions//open_auction[/interval/start]/seller/@person
XM35 //categories/category[//description/parlist/listitem]
XM36 //open_auctions/open_auction[/quantity and /current]/annotation[/description/text[/bold]/emph]/happiness
XM37 //categories/category/description/parlist[//listitem/text/emph]
XM38 //categories/category[//name]/description[/text]
XM39 //closed_auctions/closed_auction[/seller/@person]/annotation[/happiness]//description[/parlist/listitem/text//emph/bold]
XM40 //categories/category/description//parlist/listitem//text/bold
XM41 //people//person[/homepage and /@id]//profile/business
XM42 //people/person[/@id and /homepage and /creditcard]
XM43 //regions//australia/item[//name and //description//keyword]/incategory/@category
XM44 //regions[/europe]/africa/item[/mailbox/mail[//text]/from and /payment]/incategory/@category
XM45 //categories/category[/@id and //parlist/listitem/text/bold]
XM46 //categories/category[//name and /description/text//emph and /@id]
XM47 //open_auctions/open_auction[/bidder[//time and //increase]//date and //reserve]
XM48 //closed_auctions/closed_auction[/type and /quantity]//annotation/author/@person
XM49 //categories/category[/name]/description/text[/emph]
XM50 //categories/category[/@id and //description/parlist/listitem//text/bold]

E.2 SwissProt

SP01 //Entry[//MIM/@sec_id and /@class and //Keyword]
SP02 //Entry[/@id]/HSSP[/@prim_id]/@sec_id
SP03 //Entry[/Org]/Features[/TRANSMEM[/@from]/@to]
SP04 //Entry[/Features//CHAIN[/@from]/Descr]/Keyword
SP05 //Entry[/Org and /Descr]//EMBL[/@sec_id]/@prim_id
SP06 //Entry[/Org and //HSSP/@prim_id]/@class
SP07 //Entry[/@class and //PIR/@prim_id]//Keyword
SP08 //Entry[/PIR[/@prim_id]/@sec_id and //Org]
SP09 //Entry[/PIR/@sec_id]/EMBL[/@prim_id]
SP10 //Entry[/PROSITE[/@status]/@sec_id]/Org
SP11 //Entry[//PROSITE/@status and /Keyword]/EMBL[/@prim_id]
SP12 //Entry[/@class]/PIR[/@sec_id]/@prim_id
SP13 //Entry[/@mtype]/PRINTS[/@sec_id]/@prim_id
SP14 //Entry[/PIR[/@prim_id]/@sec_id and /Keyword]//Features[/TRANSMEM[//Descr]/@to]
SP15 //Entry[/PDB/@sec_id]/Features[//Descr and /CONFLICT/@from]
SP16 //Entry[/@id]/Features[/DOMAIN[/Descr]/@from]//VARSPLIC/@to
SP17 //Entry/Features[/NP_BIND]//CHAIN[/Descr]/@from
SP18 //Entry[/@mtype and /Features/BINDING/@to]
SP19 //Entry[/Org and /PFAM[/@status]/@prim_id]//Features/TRANSIT/Descr
SP20 //Entry[//REPEAT]/INTERPRO[/@sec_id and /@prim_id]
SP21 //Entry[/Mod/@Rel]/Features/CHAIN[/Descr and /@from]/@to
SP22 //Entry[//Ref]//PIR[/@prim_id]/@sec_id
SP23 //Entry[/PFAM/@status and /INTERPRO/@sec_id]//Features[//MOD_RES/@to]
SP24 //Entry[/Descr]/Features[/SITE/@from]
SP25 //Entry[/TIGR/@prim_id and /EMBL/@sec_id]
SP26 //Entry[/Species and /SWISS-2DPAGE/@prim_id]
SP27 //Entry[/INTERPRO]/Features/ACT_SITE[/@from]/Descr
SP28 //Entry[/Species]/TIGR[/@sec_id]/@prim_id
SP29 //Entry[/Org]//Features[/NON_TER/@to]/METAL//Descr
SP30 //Entry[/Features[/DISULFID[/@to]/Descr]/CHAIN]/HSSP/@prim_id
SP31 //Entry[/Features[/CHAIN]/TRANSIT/@from]
SP32 //Entry[/@mtype and //SUBTILIST[/@sec_id]/@prim_id]
SP33 //Entry[/Mod[/@type]/@date]/PRINTS/@prim_id
SP34 //Entry[/FLYBASE[/@prim_id]/@sec_id and /Gene]/@seqlen
SP35 //Entry[/Org]/Features//NP_BIND/@from
SP36 //Entry[/Features//METAL[/@to]/@from and /Species]/@class
SP37 //Entry[/@class and /EMBL]/HSSP/@sec_id
SP38 //Entry[/Org]//MGD[/@prim_id]/@sec_id
SP39 //Entry[/Org and /PDB/@sec_id]/@mtype
SP40 //Entry[/PFAM/@status and /MGD/@sec_id]
SP41 //Entry[//Keyword]//PRINTS[/@sec_id]/@prim_id
SP42 //Entry[//TIGR[/@prim_id]/@sec_id]/Org
SP43 //Entry[//Keyword and /PROSITE[/@sec_id and /@status]/@prim_id]//Features//ACT_SITE[/@to]/@from
SP44 //Entry[//INTERPRO/@sec_id]/PIR/@prim_id
SP45 //Entry[/Org]/HSSP[/@sec_id]/@prim_id
SP46 //Entry[//TUBERCULIST]/Mod[/@date and /@Rel]
SP47 //Entry[/Mod[/@type and /@Rel]/@date]/PIR/@prim_id
SP48 //Entry[//Org]//Features/METAL[/@from]/Descr
SP49 //Entry[/Keyword]/GCRDB[/@sec_id]/@prim_id
SP50 //Entry[/Mod/@date]//PRINTS[/@prim_id]/@sec_id

E.3 TreeBank

TB01 //EMPTY[//S[/_BACKQUOTES_]//VP/VBZ]//_PERIOD_
TB02 //EMPTY[//VBN and //_PERIOD_]//NP[//DT]//NN
TB03 //EMPTY[/S[//VBD and //JJ]/VP/VBN and //PP/IN]
TB04 //EMPTY[/_PERIOD_ and /S[/VBP]//TO]
TB05 //EMPTY[/S[//VP//NP//DT]/RB and //_NONE_]
TB06 //EMPTY[/S and //PP[/IN]/NP[/VBN]/NN]
TB07 //EMPTY[/S/VP/SBAR//_NONE_]/_PERIOD_
TB08 //EMPTY[//_PERIOD_ and /S/PP[/IN]/NP/NNS]
TB09 //EMPTY[//_PERIOD_]/S[//VBD]/NP[/DT]
TB10 //EMPTY[//_PERIOD_]//S[/NP/ADJP//JJ]/VB
TB11 //EMPTY[/S[/NP[/DT]/NNS and /RB]//VP]
TB12 //EMPTY/S[//NP/NNS and /VP and //RB]
TB13 //EMPTY[//_PERIOD_]//S[/NP[//CC]//NNS]//VBP
TB14 //EMPTY/S[/_PERIOD_ and /NP//NN]//ADJP/PP[//IN and //PRP_DOLLAR_]
TB15 //EMPTY//S[/NP/PRP and //_PERIOD_]
TB16 //EMPTY[/_PERIOD_]//S//VP[//VBD]//PP[//TO]//NP[//_NONE_]/CD
TB17 //EMPTY[//_PERIOD_]/_QUESTIONMARK_[/S//_NONE_]
TB18 //EMPTY[/S[//NNP]/VP[//VBD]/NP//CD]
TB19 //EMPTY//SINV[/_COMMA_ and //_PERIOD_]
TB20 //EMPTY//S//SBAR[/CC and //_NONE_]
TB21 //EMPTY[/_PERIOD_]/S[/CC and /_COMMA_]
TB22 //EMPTY[//S[//_PERIOD_ and //ADJP/PP//IN]/VP//VB and //MD]
TB23 //EMPTY[/S[/ADJP[//JJ]//RB]/VP/VBZ]
TB24 //EMPTY[//_PERIOD_]/S[/VBP]//VP//VBN
TB25 //EMPTY/S[/VP/PP//TO and /_PERIOD_]
TB26 //EMPTY//S[//VBD and /VP/PP/IN]/NP[/NN]
TB27 //EMPTY[//_PERIOD_]/S[/VBZ]//PP/NP[//JJ]/NN
TB28 //EMPTY[//_PERIOD_]/S[//SBAR/_NONE_]/VP
TB29 //EMPTY//S[//VB and //RB]/NP[//JJ and /PP/IN]//_QUOTES_
TB30 //EMPTY[/_PERIOD_]/S[/NP[/NNS]//NNP]//ADJP[/RB]
TB31 //EMPTY/S[/_PERIOD_]//NP[//NNP and //_LRB_]
TB32 //EMPTY[//VP//ADJP/JJ]//S//_PERIOD_
TB33 //EMPTY[/_PERIOD_]/SBAR[//IN]//S[//NNP and //VP[/VBZ]//PP//NN]
TB34 //EMPTY[/_PERIOD_]//S//VP[/VBZ]/PP[/IN]/NP//ADJP/JJ
TB35 //EMPTY[/_PERIOD_]/S[//VP[/VBN]//TO]/NP[//PP/IN]/NNPS
TB36 //EMPTY[/S//VP[//VB]/PP[//VBG]/TO]
TB37 //EMPTY[/_PERIOD_]/S[/VP[/VBZ]/NP and //_COMMA_]
TB38 //EMPTY[//_PERIOD_]/S[/VP[/NP/NN]/VBD and //SBAR]//VP-1/VBG
TB39 //EMPTY[/_PERIOD_]/S/NP[/CD and /NNS]
TB40 //EMPTY[//_PERIOD_ and /S[/NP[//JJ]/NN]/RB]
TB41 //EMPTY[/S[//VP//VB and /VBZ]/RB and /_PERIOD_]
TB42 //EMPTY[/_PERIOD_]//S/SINV/VP[//VBZ]
TB43 //EMPTY[/_PERIOD_]/S[/NP/PRP]/PP//IN
TB44 //EMPTY[/S[/VP[/PP[/NP/NN]//IN]/RB]//_PERIOD_]
TB45 //EMPTY[//_PERIOD_]//S/SBARQ[/SQ]/_NONE_
TB46 //EMPTY//S[//_BACKQUOTES_]/NP[/PRP]
TB47 //EMPTY[//X[//_COLON_]//NP//NNPS]/_PERIOD_
TB48 //EMPTY[/_PERIOD_ and /S[/PP[/NP/NNP]//CC]/_COMMA_]
TB49 //EMPTY[/_PERIOD_]/S/SINV/VP[//VBD]
TB50 //EMPTY[/_PERIOD_]//S[/VP[//VBD]//PP/IN]/NP[//NNS]/DT

F. QUERIES FOR PERFORMANCE COMPARISON

F.1 XMark

1
```
for $a in //categories/category[.//description[.//text] and ./@id]
return ($a)
```

2
```
for $a in //closed_auctions/closed_auction[.//price and ./buyer[./@person]]
return ($a)
```

3
```
for $a in //regions/samerica/item[./mailbox[./mail[./to]]]/incategory/@category
return ($a)
```

4
```
for $a in //open_auctions//open_auction[./reserve and .//bidder[./personref[./@person]]]/itemref[./@item]
return ($a)
```

5
```
for $a in //regions/namerica/item[.//from]/incategory/@category
return ($a)
```

6
```
for $a in //regions
for $b in $a/namerica
for $c in $b/item
for $d in $c//from
for $e in $c/incategory
for $f in $e/@category
return ($a, $b, $c, $d, $e, $f)
```

7
```
for $a in //regions
for $b in $a//europe
for $c in $a/namerica
for $d in $c/item
for $e in $d/description
for $f in $e/parlist
for $g in $e//listitem
for $h in $g/text
return ($a, $b, $c, $d, $e, $f, $g, $h)
```

8
```
for $a in //people
for $b in $a/person
for $c in $b/emailaddress
for $d in $b//creditcard
for $e in $b/@id
return ($a, $b, $c, $d, $e)
```

9
```
for $a in //closed_auctions
for $b in $a/closed_auction
for $c in $b/type
for $d in $b/quantity
return ($a, $b, $c, $d)
```

10
```
for $a in //people
for $b in $a/person
for $c in $b/profile
for $d in $c//business
for $e in $c/interest
for $f in $e/@category
return ($a, $b, $c, $d, $e, $f)
```

F.2 TreeBank

1
```
for $a in //EMPTY[.//S[./_BACKQUOTES_ and .//VP[./VBZ]]]//_PERIOD_
return ($a)
```

2
```
for $a in //EMPTY[.//PP[./IN]]/S[.//VBD and .//JJ]/VP[./VBN]
return ($a)
```

3
```
for $a in //EMPTY[./_PERIOD_ and ./S[./VBP and .//TO]]
return ($a)
```

4
```
for $a in //EMPTY[./S]//PP[./IN]/NP[./VBN and ./NN]
return ($a)
```

5
```
for $a in //EMPTY[.//_PERIOD_]/S[./PP[./IN and ./NP[./NNS]]]
return ($a)
```

6
```
for $a in //EMPTY
for $b in $a/S
for $c in $b//NP
for $d in $c/NNS
for $e in $b/VP
for $f in $b//RB
return ($a, $b, $c, $d, $e, $f)
```

7
```
for $a in //EMPTY
for $b in $a//_PERIOD_
for $c in $a//S
for $d in $c/NP
for $e in $d//CC
for $f in $d//NNS
for $g in $c//VBP
return ($a, $b, $c, $d, $e, $f, $g)
```

8
```
for $a in //EMPTY
for $b in $a//S
for $c in $b/NP
for $d in $c/PRP
for $e in $b//_PERIOD_
return ($a, $b, $c, $d, $e)
```

9
```
for $a in //EMPTY
for $b in $a/_PERIOD_
for $c in $a//S
for $d in $c//VP
for $e in $d//VBD
for $f in $d//PP
for $g in $f//TO
for $h in $f//NP
for $i in $h//_NONE_
for $j in $h/CD
return ($a, $b, $c, $d, $e, $f, $g, $h, $i, $j)
```

10
```
for $a in //EMPTY
for $b in $a/S
for $c in $b//NNP
for $d in $b/VP
for $e in $d//VBD
for $f in $d/NP
for $g in $f//CD
return ($a, $b, $c, $d, $e, $f, $g)
```